# z6 – The Next-Generation Mainframe Microprocessor

Charles Webb
IBM Systems &Technology Group, Development
IBM Fellow



IBM Systems

## z6 Highlights

- **New high-frequency z/Architecture microprocessor core**
  - **>4 GHz operation in system**
- **4 cores per die**
  - **Each with 3MB private 2nd-level cache**
- **Accelerator engines**
  - **Data compression**
  - **Cryptographic functions**
  - **Decimal floating point**
- **Integrated SMP communications**
  - **Switch connects cores to SMP Hub chip**
    - **Shared cache and SMP fabric**
  - **Memory bus controller**
  - **I/O bus controller**
  - **EI3 technology up to 3 GHz bus speeds**
- **System interfaces**
  - **2 x 48 GB/s SMP Hub**
  - **4 x 13 GB/s Memory**
  - **2 x 17 GB/s I/O**

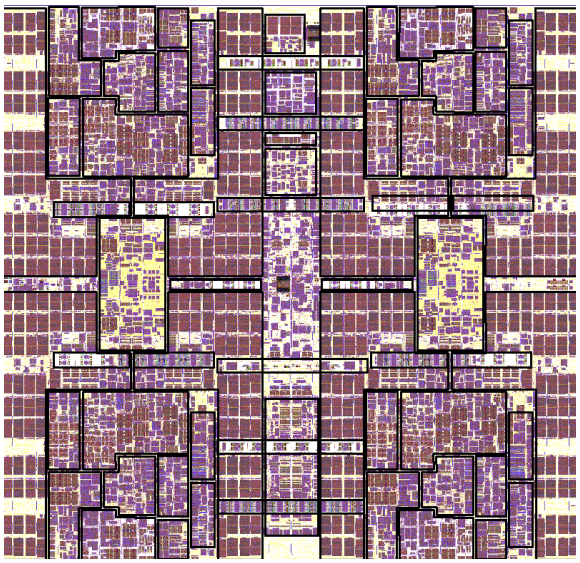

- **991M Transistors**
- **138 Mb SRAM**
- **6 km wire**
- **21.7 X 20.0 mm die**
- **1188 signal / 8765 total chip I/Os**

# Relationship to Power6

- **Siblings, not identical twins**
- **Share lots of DNA**
  - **–IBM 65nm SOI technology**
  - **–Design building blocks:**
    - **•latches, SRAMs, regfiles, dataflow elements**
  - **–Large portions of FXU, BFU, DFU, MC, GX**
  - **–EI3 interface technology**
  - **–Core pipeline design style**
    - **•High-frequency, low-latency, mostly-in-order**
  - **–Many designers**
- **Different personalities**
  - **–Very different ISAs => very different cores**
  - **–Cache hierarchy and coherency model**
  - **–SMP topology and protocol**
  - **–Chip organization**
  - **–z6 optimized for Enterprise Data Serving Hub**

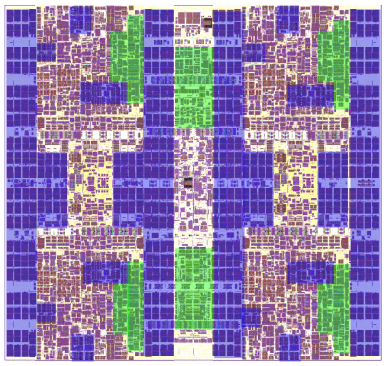

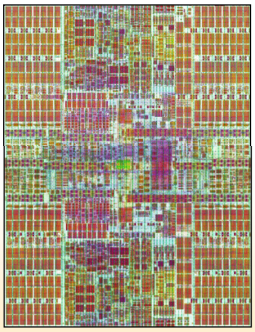





# z6 Instruction Set Architecture

- **Continues line of upward-compatible mainframe processors**
  - **Application compatibility since 1964**
  - **Supports all z/Architecture-compliant OSes**

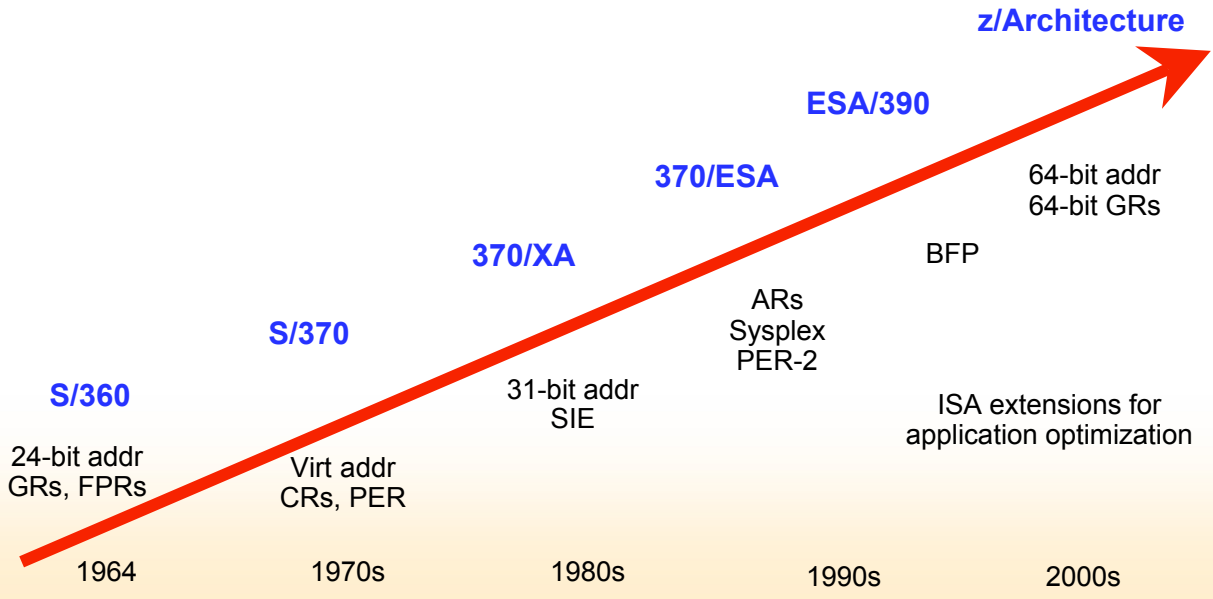

# z6 Architecture

- Continues line of upward-compatible mainframe processors

- **Rich CISC ISA**
  - **894 instructions (668 implemented entirely in hardware)**
  - **24, 31, and 64-bit addressing modes**
  - **Multiple address spaces robust inter-process security**
  - **Multiple arithmetic formats**
  - **Industry-leading virtualization support**
    - **High-performance logical partitioning via PR/SM**
    - **Fine-grained virtualization via zVM scales to 1000s of images**
  - **Precise, model-independent definition of hardware / software interface**

- **Architectural extensions for z6**
  - **50+ instructions added to improve compiled code efficiency**
  - **Enablement for software/hardware cache optimization**
  - **Support for 1MB page frames**
  - **Full hardware support for Decimal Floating Point**




# z6 Continues the CMOS Mainframe Heritage

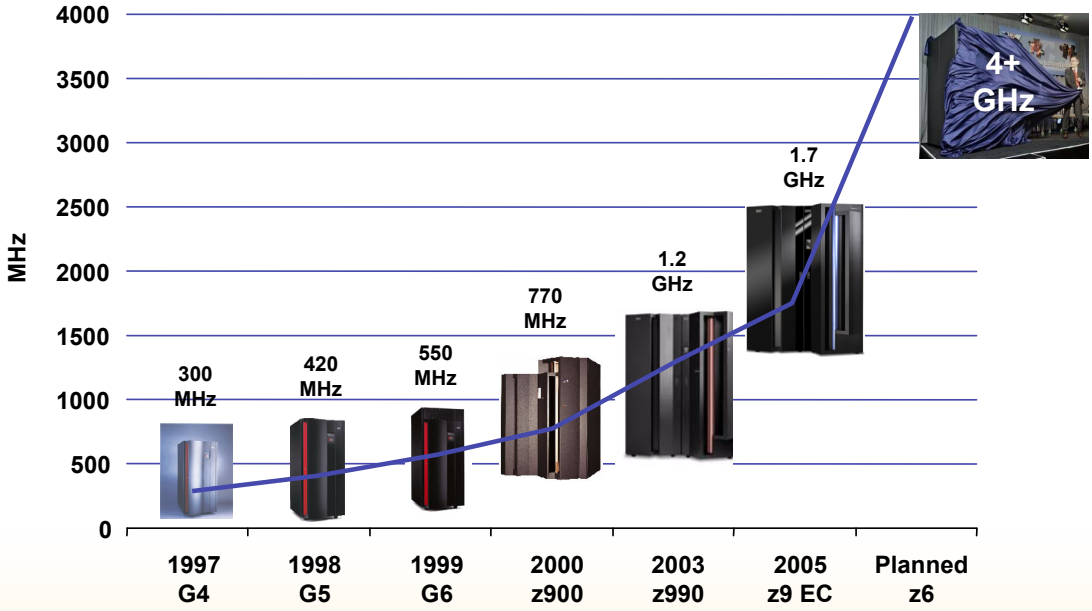


- **G4 – 1st full-custom CMOS S/390**
- **G5 – IEEE-standard BFP; branch target prediction**
- **G6 – Cu BEOL**
- **z900 – Full 64-bit z/Architecture**
- **z990 – Superscalar CISC pipeline**
- **z9 EC – System level scaling**

# z6 Microprocessor Core

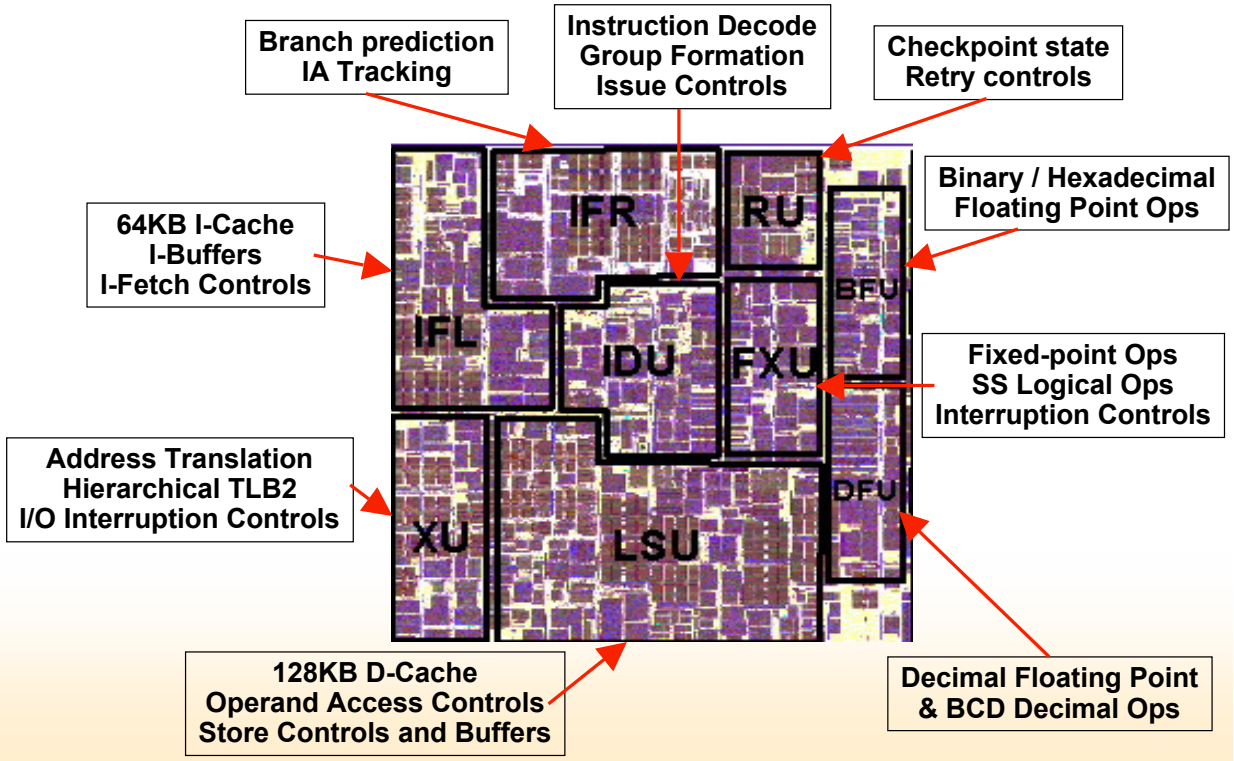





# High-Frequency Design

- **Frequency-optimized methodology**
  - **Logic / circuit co-design**
  - **Reduced latch delay overhead**
  - **Arithmetic dataflow stacks shared with Power6**
- **Minimize critical pipeline latencies**
- **Isolate hardware complexity of CISC ISA**
- **Optimization of cache lookup path**
- **Aggressive branch prediction**

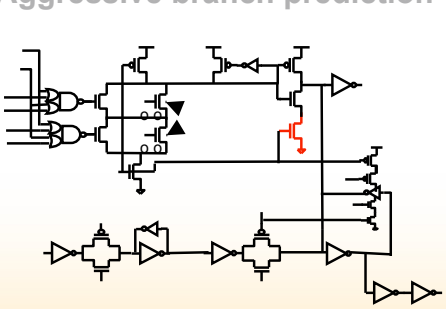

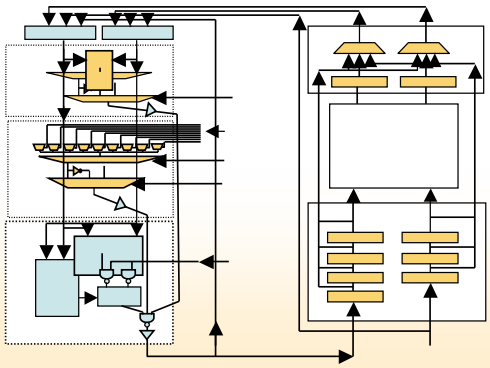

# High-Frequency Design

- Frequency-optimized methodology
- **Minimize critical pipeline latencies**
  - **Overall pipeline depth increased significantly from predecessor**
    - **15FO4 vs. 28FO4 cycle**
  - **Minimal increase in performance-critical core loop latency**
  - **ISA complexity pushed to front end of pipeline**
    - **outside of key dependency loops**
  - **Non-stalling pipeline:**
    - **once issued, instruction either finishes or is recycled**
- Isolate hardware complexity of CISC ISA
- Optimization of cache lookup path
- Aggressive branch prediction





# z6 Core Pipeline

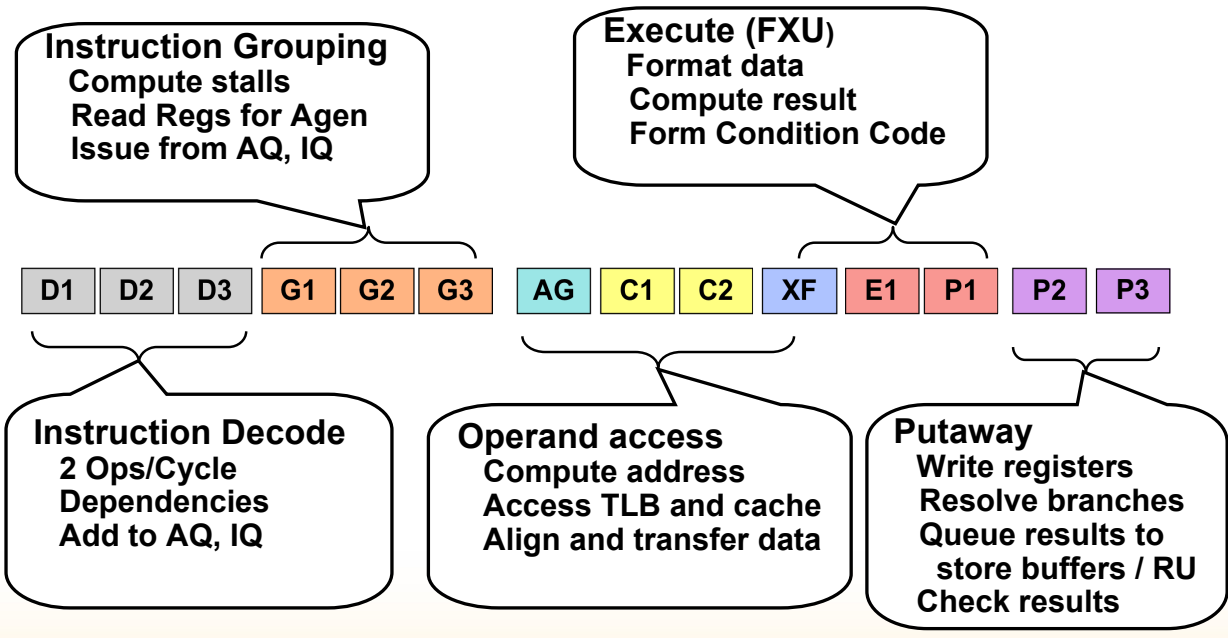

# z6 Core Pipeline

**15 FO4 / cycle**

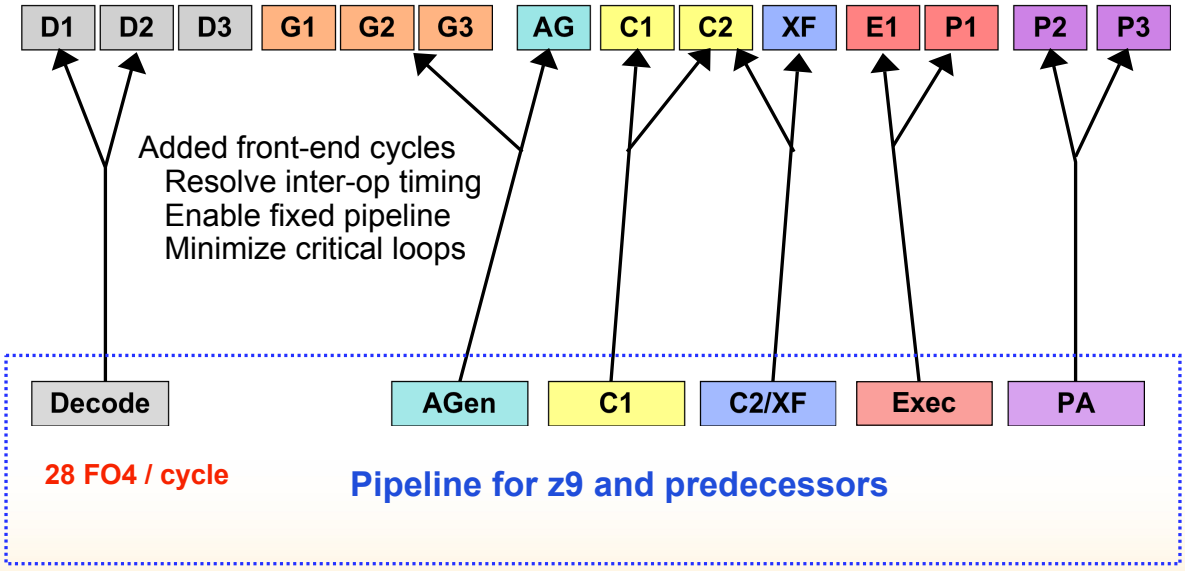


**28 FO4 / cycle**

**Pipeline for z9 and predecessors**



# z6 Core Pipeline

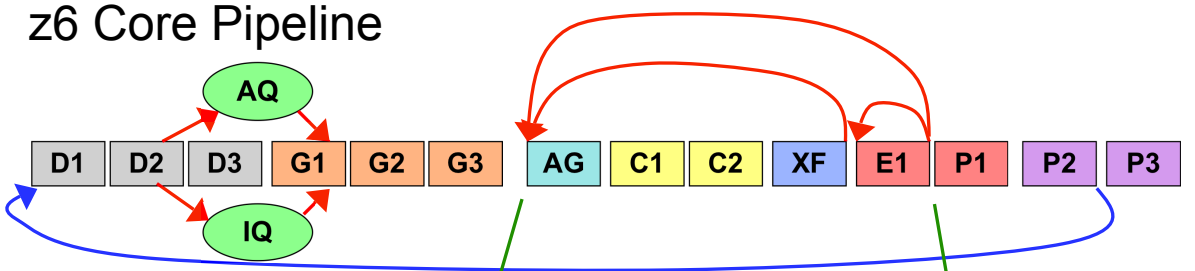


**Performance-critical heart of pipeline**

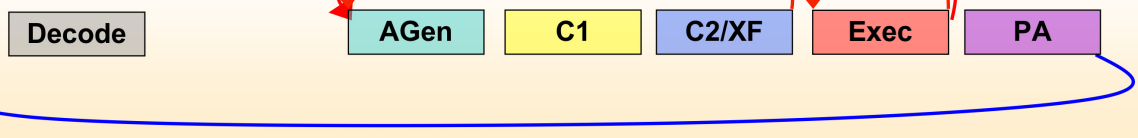

# z6 Core Pipeline

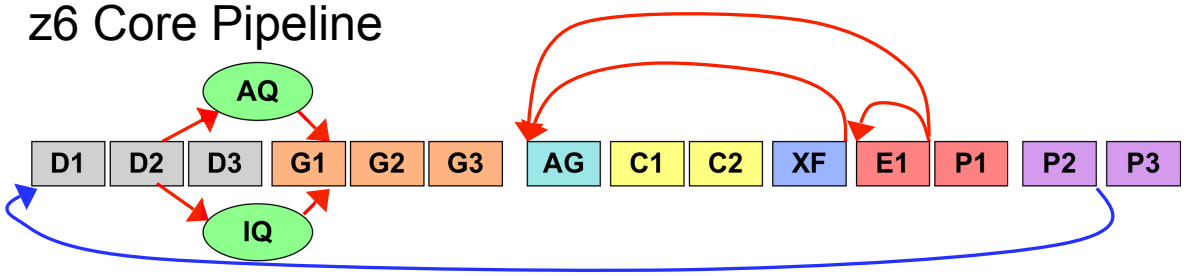


| Pipeline Loop | Cycles | | FO4 | |
|---|---|---|---|---|
| | z9 | z6 | z9 | z6 |
| Load-FXU | 0 | 0 | 0 | 0 |
| FXU-FXU | 1 | 1 | 28 | 15 |
| Load-Load | 3 | 4 | 84 | 60 |
| FXU-Load | 4 | 5 | 112 | 75 |
| Branch wrong | 6+ | 13+ | 168+ | 195+ |

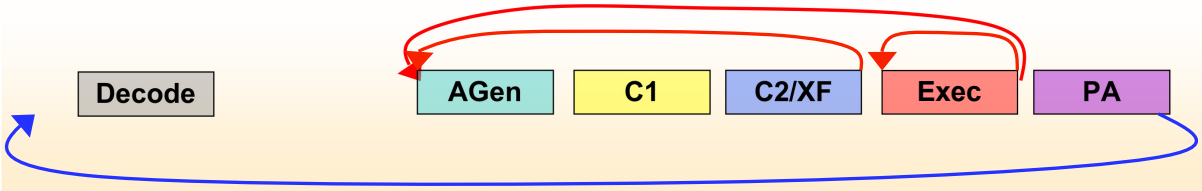





# High-Frequency Design

- Frequency-optimized methodology
- Minimize critical pipeline latencies
- **Isolate hardware complexity of CISC ISA**
  - **Multi-pass handling of special cases**
  - **Leverage millicode for complex operations**
- **Intense optimization of Address_Add / Cache lookup / FXU path**
  - **Cache set prediction; latch boundary adjustment**
  - **Performance / array / logic co-optimization**

    **➔128KB D-Cache and 4-cycle loop**
- Aggressive branch prediction

# Operand Access Path

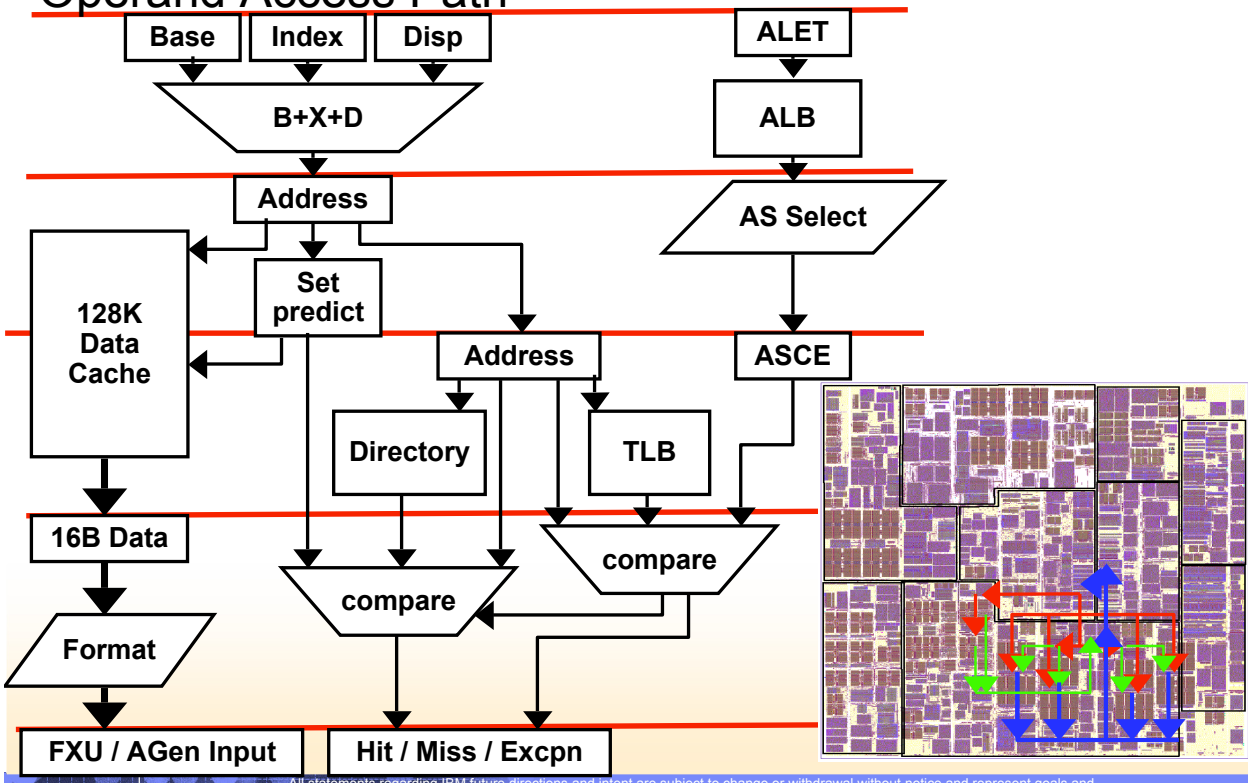





# High-Frequency Design

- Frequency-optimized methodology
- Minimize critical pipeline latencies
- Isolate hardware complexity of CISC ISA
- Optimization of cache lookup path
- **Aggressive branch prediction**
  - **Minimize impact of long front end of pipe**

# z6 Branch Prediction

- **Multiple history-based prediction mechanisms**
  - 2 level Branch Target Buffer
  - Filtered Pattern History Table
  - Tagged multi-target prediction
  - Level 2 BTB data compression

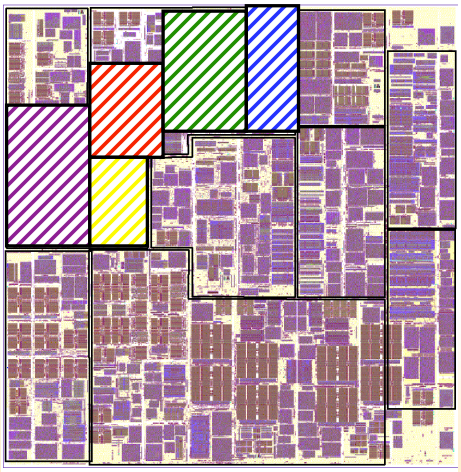

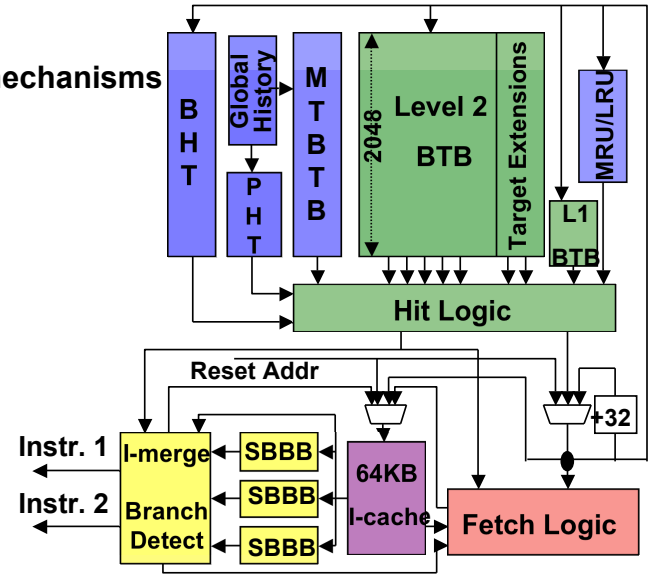




# z6 Branch Prediction

- **Multiple history-based prediction mechanisms**
  - 2 level Branch Target Buffer
  - Filtered Pattern History Table
  - Tagged multi-target prediction
  - Level 2 BTB data compression

- **BTB searches run ahead of I-fetch**
  - Enables I-cache prefetching of targets
  - Prevents fall-through false misses in I-cache
- **Tight interaction with I-Fetch and IDU**
  - Basic block merging
  - Relative branch snooping
- **z/Architecture Optimizations**
  - Millicode entry / exit prediction
  - EXecute instruction optimization
  - Efficient handling of "indirect" branches

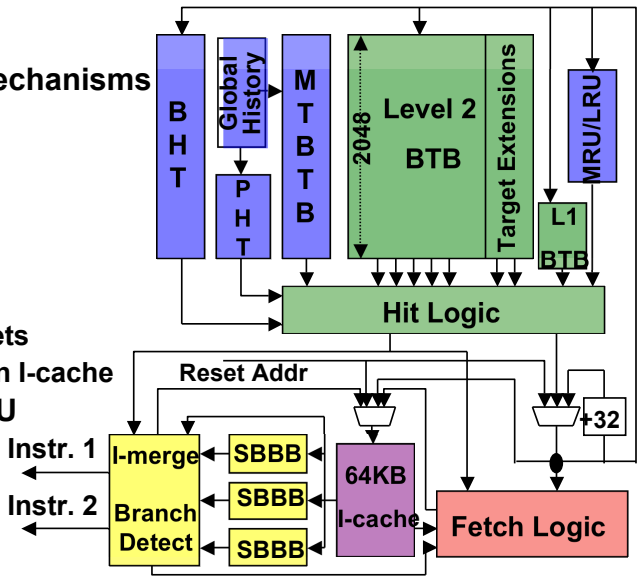


© 2007 IBM Corporation
All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.
IBM Systems

# Decimal Floating Point Accelerator

- **Meets requirements of business and human-centric applications**
  - ☞**Performance, Precision, Function**
    - **Avoids rounding and other problems with binary/decimal conversions**
    - **Improved numeric functionality over legacy BCD operations**
    - **Much of commercial computing is dominated by decimal data and decimal operations**

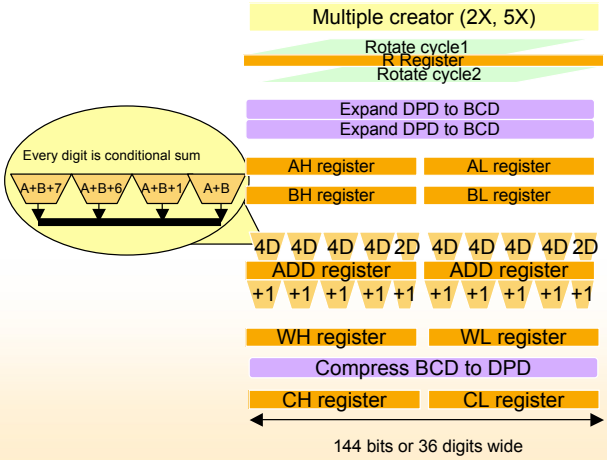





# Decimal Floating Point Accelerator

- **Meets requirements of business human-centric applications**
  - ☞**Performance, precision, Function**

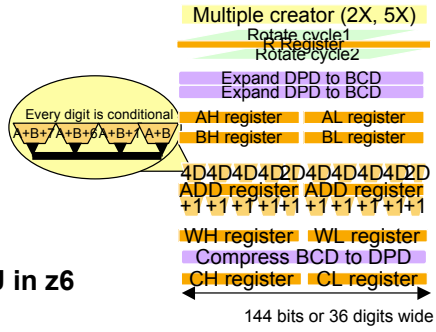


- **z6 DFU co-developed with Power6**
  - –**Common architecture operations and semantics**
  - –**Common dataflow elements**
  - –**Mainframe legacy BCD operations mapped onto DFU in z6**
- **Growing industry support for DFP standardization**
  - –**Java BigDecimal,C#, XML, C/C++, GCC, COBOL all supporting IEEE 754R**
  - –**Endorsed by key software vendors including Microsoft and SAP**
  - –**Open standard definition led by IBM**

# Compression and Cryptography Accelerator

- **Data compression engine**
- **Cryptography engine**
- **Accelerator unit shared by 2 cores**

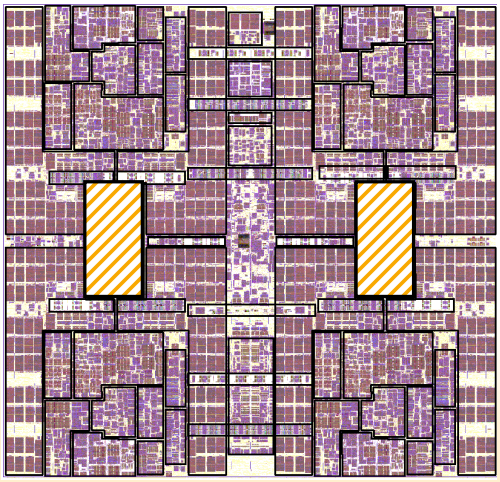

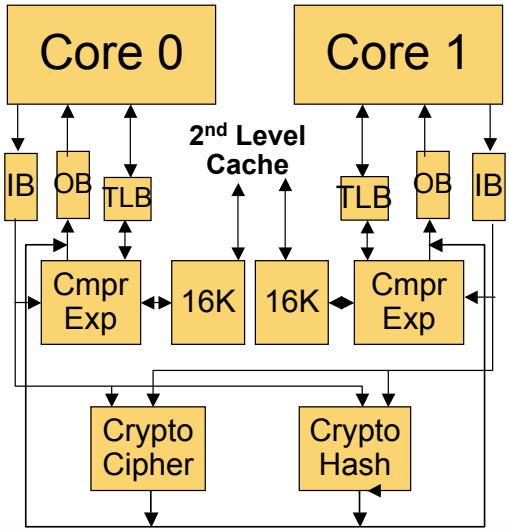




# Compression and Cryptography Accelerator

- **Data compression engine**
  - **Static dictionary compression and expansion**
  - **Dictionary size up to 64KB (8K entries)**
    - **Local 16KB caches for dictionary data**
  - **Up to 8.8 GB/sec expansion**
  - **Up to 240 MB/sec compression**
- **Cryptography engine**
  - **DES (DEA, TDEA2, TDEA3)**
  - **SHA-1 (160 bit)**
  - **SHA-2 (256, 384, 512 bit)**
  - **AES (128, 192, 256 bit)**
  - **290-960 MB/sec bulk encryption rate**
- **Accelerator unit shared by 2 cores**
  - **Independent compression engines**
  - **Shared cryptography engines**
  - **Co-operates with core millicode**
  - **Direct path into core store buffers**

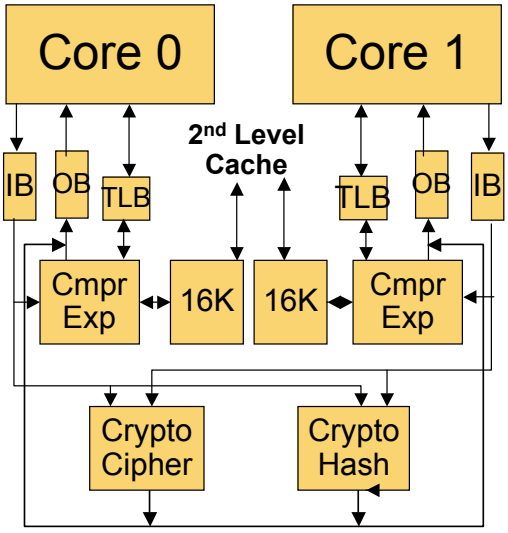


© 2007 IBM Corporation
All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.
IBM Systems

# Industry-Leading Error Detection and Recovery

- **Fine-grained redundancy and checking throughout design**
  - **–ECC on 2nd and 3rd-level caches, store buffers, R-Unit state array**
  - **–Parity on all other arrays, register files**
  - **–Parity or residue checking on data/address/execution flow**
  - **–Extensive functional, parity, and state checking on control logic**
  - **–Over 20,000 error checkers in chip**
- **Full architected state of processor buffered in R-Unit with ECC**
  - **–Allows precise core retry for almost all hardware errors**
  - **–Dynamic, transparent core sparing in the event of hard error in core**
- **Machine check architecture allows precise software recovery**
  - **–Minimizes system impact in rare case of unrecoverable failure**

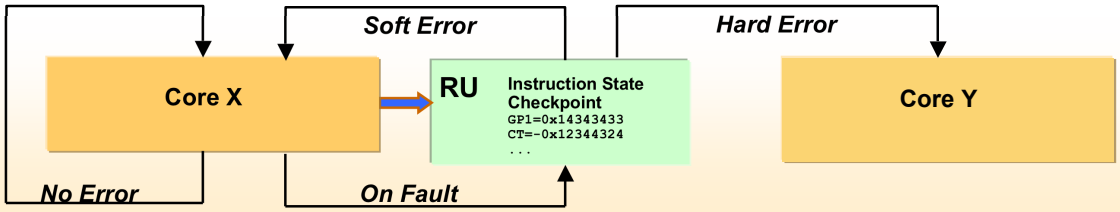





# Industry-Leading Error Detection and Recovery

- Fine-grained redundancy and checking throughout design
  - –ECC on 2nd and 3rd-level caches, store buffers, R-Unit state array
  - –Parity on all other arrays, register files
  - –Parity or residue checking on data/address/execution flow
  - –Extensive functional, parity, and state checking on control logic
  - –Over 20,000 error checkers in chip
- Full architected state of processor buffered in R-Unit with ECC
  - –Allows precise core retry for almost all hardware errors
  - –Dynamic, transparent core sparing in the event of hard error in core
- Machine check architecture allows precise software recovery
  - –Minimizes system impact in rare case of unrecoverable failure

☞ **Bullet-proof computing:**

✓ **Never lose data**

✓ **Never go down**

© 2007 IBM Corporation
All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.
IBM Systems

# SMP Hub Chip

- **Connects multiple z6 Processor chips**
  - –**48 GB/Sec bandwidth per processor**
- **Shared 3rd-Level cache**
  - –**24MB SRAM Cache**
  - –**Extended directory**
    - • **Partial-inclusive discipline**
  - –**Hub chips can be paired**
    - • **48MB shared cache**
- **Low-latency SMP coherence fabric**
  - –**Robust SMP scaling**
  - –**Strongly-ordered architecture**
- **Multiple hub chips/pairs allow further SMP scaling**

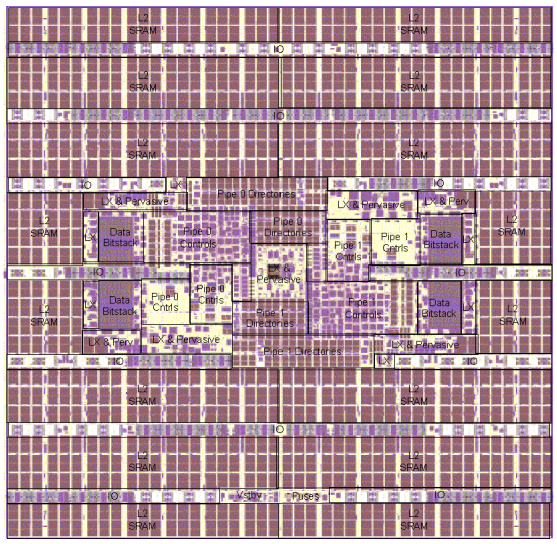


- **1.6 Billion Transistors**
- **242 Mb SRAM**
- **3 km wire**
- **20.8 X 21.4 mm die**
- **2419 signal / 7984 total chip I/Os**




# Energy Efficiency

- **Mainframe focus on system and data center efficiency**
  - – **Consolidation of many servers onto one system**
  - ☞ **Consistent performance at sustained high utilization**
    - • **Resilience in face of changing workloads**
  - – **Leverages virtualization capabilities of PR/SM, zVM, zOS**
- **System z designs are optimized for scale-up data serving**
  - – **SMP Hub design enables robust scaling across wide spectrum of workloads**
  - – **Centralized SMP fabric minimizes fabric logic per core**
    - • **Extended on z6 via 4-core processor chip**
  - – **MRU cooling allows dense package and reduces leakage power**
  - – **Extensive hardware support for multi-level virtualization**
- **Chip-level power optimization applied to z6 design**
  - – **Local clock gating to limit maximum dynamic power**
  - – **Millicode sleep mode for wait/spare/stop states**

# Conclusion

- **z6 processor takes mainframe computing to the next level**
  - **New high-frequency 4-core microprocessor**
  - **Advanced pipeline design**
- **z6 is specifically designed and optimized for mainframe systems**
  - **Full z/Architecture compatibility**
  - **New features enhance enterprise data serving performance**
  - **Industry-leading virtualization capabilities**
  - **Energy efficiency at system and data center levels**
  - **Extends heritage of industry-leading RAS**





*Thank You!*

Charles Webb
cfw@us.ibm.com

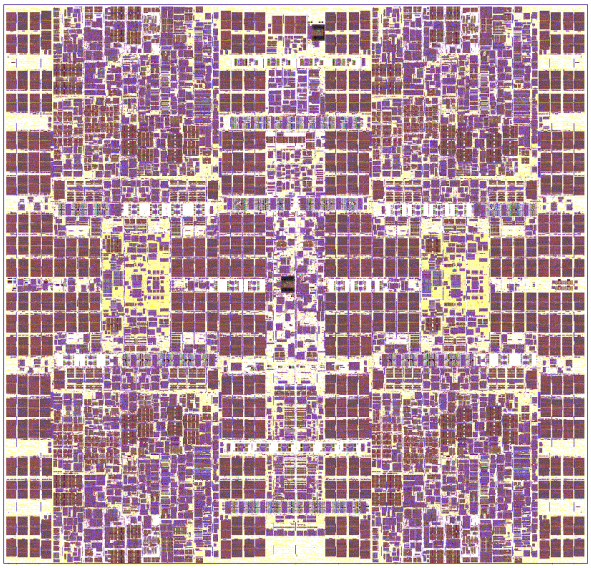

© 2007 IBM Corporation
All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.